Ethical and Bias-Aware Machine Learning for Mental Health and Behavioral Analytics

Adeel Farooq
Department of Computer Science,
Bahria University Islamabad Campus, Pakistan

Sana Khalid
Department of Computer Science,
Bahria University Islamabad Campus, Pakistan

Hamza Rauf
Department of Computer Science,
Bahria University Islamabad Campus, Pakistan

Submitted on: August 14, 2021
Accepted on: September 22, 2021
Published on: November 10, 2021

DOI: 10.5281/zenodo.18065535

Abstract—**Machine learning systems are increasingly used to infer mental health conditions, emotional states, and behavioral patterns from digital traces such as text, speech, physiological signals, and interaction logs. While these systems promise scalable and early mental health insights, they also introduce significant ethical risks related to bias, privacy, interpretability, and potential harm. This article investigates ethical and bias-aware machine learning frameworks for mental health and behavioral analytics. We analyze common sources of bias across data, model design, and deployment contexts, and examine how they affect fairness and reliability in mental health inference. Building on recent advances in deep learning, natural language processing, and affective computing, we propose a multi-layered ethical machine learning architecture that integrates bias detection, fairness constraints, and explainability mechanisms. Empirical evaluations using representative behavioral datasets demonstrate that incorporating ethical controls improves robustness and reduces disparity across demographic groups while maintaining predictive performance. The findings highlight the necessity of embedding ethical considerations directly into the machine learning lifecycle for responsible mental health analytics.**

Index Terms—**Ethical AI, bias-aware learning, mental health analytics, behavioral modeling, fairness, explainable machine learning**

I. Introduction

The growing availability of digital behavioral data has enabled machine learning models to infer psychological states such as depression, stress, emotional valence, and cognitive engagement. Social media posts, mobile sensing data, wearable signals, and online interaction logs are increasingly analyzed to support mental health screening and behavioral insights. Deep learning models have shown strong performance in extracting complex patterns from such data, including sequential and multi-modal representations [1]–[3].

Despite these advances, mental health analytics presents unique ethical challenges. Behavioral data is deeply personal, context-sensitive, and often reflects vulnerable populations. Bias introduced through unbalanced datasets, cultural assumptions, or opaque models can lead to misclassification and unfair outcomes. For example, linguistic markers of depression may vary significantly across demographics, languages, and socio-economic contexts [4], [5]. Without bias-aware mechanisms, machine learning systems risk reinforcing existing inequalities.

This article addresses the need for ethical and bias-aware machine learning in mental health and behavioral analytics. The contributions are threefold. First, we provide a structured review of relevant machine learning approaches and ethical risks. Second, we propose an integrated architecture that embeds fairness, transparency, and accountability into the modeling pipeline. Third, we empirically evaluate the impact of ethical constraints on model performance and fairness using representative behavioral datasets.

II. Literature Review

Ethical and bias-aware machine learning for mental health and behavioral analytics sits at the intersection of predictive modeling, human centered decision support, and responsible data practice. Prior work across behavioral prediction, affect recognition, and healthcare analytics provides both technical building blocks and cautionary lessons about reliability, generalization, and unintended harm. This review groups the literature into thematic areas that inform the design choices in bias-aware mental health systems.

A. Digital behavioral signals for mental health inference

A substantial body of work demonstrates that mental health indicators can be inferred from digital traces such as online text, interaction patterns, and temporal activity signals. Predictive approaches on behavioral data often rely on sequence modeling to capture how symptoms or signals evolve over time. Recurrent and gated architectures have been used to model temporal dependencies for classification tasks where the target signal is sparse and noisy [2], [6]. Related studies show that performance gains often come from careful representation learning and regularization rather than model depth alone, especially when data distributions shift across cohorts [7]–[9].

Work on prediction tasks with limited labels highlights a recurring risk: improvements on aggregate metrics can hide systematic errors on minority subgroups. This is common in settings where training samples over-represent particular demographics or behavioral contexts. As a result, the same modeling choices that boost accuracy can increase disparity if the learning objective does not account for group wise reliability [10], [11]. These findings motivate fairness constraints and model diagnostics tailored to vulnerable populations.

B. Text, sentiment, and personality signals

Text driven mental health analytics frequently uses sentiment, emotion, and linguistic markers. Sentiment analysis pipelines have matured and are often treated as general purpose components, yet subtle differences in language use can produce biased inferences across communities. Studies on sentiment modeling emphasize feature learning that captures context and pragmatic cues rather than relying on shallow word statistics [5]. Predictive modeling for personality and behavioral traits similarly shows that language features can be informative, but also confounded by culture, education, and platform norms [4].

Social media based prediction work illustrates that model performance depends on domain adaptation and temporal drift handling, because the meaning of cues changes with topic, platform moderation, and community conventions [1], [2]. These results reinforce an ethical requirement: systems should quantify uncertainty and avoid overconfident outputs, especially when deployed for screening or triage. Robust learning methods that stabilize representations under distribution change become directly relevant to bias-aware mental health inference [8], [10].

C. Physiological and multimodal affect recognition

Mental health and behavioral analytics increasingly leverages physiological signals and multimodal inputs such as EEG, audio, and visual cues. Emotion recognition from brain signals demonstrates that deep architectures can learn discriminative patterns, but these models are sensitive to noise, sensor placement, and participant variability [3]. Multimodal approaches that fuse visual and acoustic streams further increase predictive power, yet introduce complex failure modes when one modality is missing or lower quality for a subset of users [12].

Sequence aware fusion methods, including CNN LSTM and hybrid ResNet LSTM designs, have been reported to improve temporal consistency and capture cross time dynamics [13]. However, multimodal systems can amplify bias through unequal sensor access, different recording environments, or device quality. In practice, this means fairness assessment cannot be limited to labels and demographics only. It must also include measurement conditions and data acquisition pathways, which are often correlated with socio-economic status and accessibility.

D. Modeling patterns over time, structure, and context

Behavioral signals are structured. They appear as sequences, graphs, and high dimensional embeddings rather than independent samples. Work on deep sequential modeling for complex prediction tasks shows that temporal encoders can capture long range dependencies but may overfit to frequent patterns and ignore rare yet clinically meaningful events [13]. Graph based learning further provides a way to represent relational context, such as social interactions or co-occurring behavioral markers [14]. In mental health settings, graph representations are useful for modeling support networks, discussion dynamics, or symptom co-occurrence. At the same time, relational modeling can leak sensitive information through neighborhood effects, raising fairness and privacy concerns if group membership is implicitly encoded.

Context aware learning approaches highlight that the same observed behavior can have different meanings depending on environment and baseline norms. This is one of the reasons ethical evaluation needs scenario based testing rather than a single test set split. Studies that emphasize domain generalization and reliability under variation provide practical guidance for mental health deployment, where usage environments are rarely controlled [8], [10].

E. Bias sources, data imbalance, and mitigation strategies

Bias often starts in the dataset. Imbalanced class distributions are common in mental health data, where positive cases may be rare and labels can be noisy. Synthetic oversampling and balancing techniques have been widely used to address class imbalance [15]. Late fusion and feature selection strategies can also improve classification stability by reducing noise and redundancy [16]. While these techniques help, they do not automatically guarantee fairness. In fact, oversampling can increase false positives for certain groups if the synthetic generation process mirrors majority group characteristics.

Representation learning frameworks that explicitly separate shared structure from group specific artifacts are increasingly relevant. Semi supervised and self training methods aim to leverage unlabeled data without reinforcing bias, but they require careful calibration to avoid confirmation loops [11]. In mental health analytics, where ground truth is hard to obtain and self reports vary, these risks are heightened. The literature suggests that bias-aware objectives should be paired with monitoring tools that track error gaps, calibration gaps, and stability across subpopulations [8], [10].

F. Robustness, adversarial risk, and safety

Robustness is an ethical issue in mental health applications because unstable predictions can cause harm through unnecessary alarm or missed support. Robust learning work emphasizes resilience to noise, perturbations, and distribution shift, which are common in real behavioral data [10]. Adversarial learning studies show that models can be manipulated or can fail under small input changes, a risk that matters for text based screening and conversational systems where prompts can be crafted intentionally or unintentionally [17]. Even when there is no attacker, natural variation in language and behavior can act like adversarial perturbations.

These findings imply that ethical system design should include stress testing. Instead of reporting only average accuracy, evaluation should include worst case slices and robustness metrics that reflect practical usage. Work on robust architectures and stability oriented training supports the inclusion of safety checks and uncertainty estimation as first class components in mental health pipelines [8], [10].

G. Explainability, surrogate models, and human oversight

Explainability is central to ethical mental health analytics because stakeholders must understand why a model flags risk, especially when interventions can affect employment, insurance, or clinical decisions. Research on surrogate modeling and interpretable approximations provides techniques for creating human readable explanations while preserving predictive performance [18]. Algorithm selection and meta learning perspectives emphasize that model choice should be treated as a decision process shaped by constraints and stakeholder needs, not only by raw performance [19].

In decision support contexts, agent based and knowledge driven methods provide a complementary lens: they emphasize traceability, rule based reasoning, and controllable decision logic [20]. This is particularly relevant for mental health triage, where a system should support clinicians or counselors rather than replace judgment. Novel decision frameworks and structured DSS approaches further underline the importance of transparent trade offs and auditable decision pathways [21]. Taken together, the literature supports hybrid designs where machine learning provides predictive signals and explanation layers translate them into actionable, reviewable evidence.

H. Healthcare and clinical analytics parallels for fairness and ethics

Although mental health analytics has unique sensitivities, adjacent work in clinical imaging and physiological diagnostics provides transferable lessons about bias, validation, and safe deployment. Automated diagnostic systems in medical imaging demonstrate that high headline performance can mask failure on underrepresented patient groups or uncommon presentations [22], [23]. ECG focused deep models also illustrate that performance depends on signal quality and cohort variability, reinforcing the need for stratified evaluation and careful preprocessing [24]. In maternal and fetal health modeling, predictive tools must respect clinical context and uncertainty, since the cost of errors is high and labels can be imperfect [25].

These healthcare studies highlight practical governance practices that mental health systems can adopt: data documentation, subgroup audits, and post deployment monitoring. They also reinforce an ethical stance: model outputs should be treated as probabilistic support signals rather than definitive diagnoses.

I. Behavioral analytics beyond health and implications for responsible use

Behavioral analytics also appears in domains such as learning platforms and user engagement modeling. Studies on large scale learning data show that predictive models can profile users and influence opportunities, which raises fairness and transparency concerns similar to those in mental health [26]. Exploratory analyses and domain specific modeling work demonstrate that behavioral labels are often proxies and can encode institutional bias [27]. Methods that enhance feature extraction and representation learning can improve performance, but also make systems harder to interpret, increasing the need for explanation and governance [28], [29].

Recent work on deep learning in applied settings emphasizes that operational factors such as latency, data pipelines, and feedback loops can change model behavior after deployment [30], [31]. For mental health, this means ethical design must consider the full lifecycle, including data refresh practices, drift detection, and controlled retraining to prevent silent performance decay.

J. Synthesis and gaps

Across these themes, a consistent message emerges. High performing behavioral models are feasible, but ethical deployment requires more than tuning hyperparameters. The literature supports three design requirements: (i) fairness evaluation that goes beyond aggregate metrics and checks subgroup stability [10], [15]; (ii) robustness testing that treats distribution shift and perturbations as expected conditions [8], [17]; and (iii) explainability and oversight mechanisms that translate predictions into reviewable evidence and enable accountability [18]–[20]. These gaps motivate an integrated ethical learning pipeline for mental health and behavioral analytics, where bias controls, interpretability, and monitoring are embedded into the modeling workflow rather than added after deployment.

III. METHODOLOGY

This section presents the proposed ethical and bias-aware machine learning framework for mental health and behavioral analytics.

A. Ethical Learning Architecture

The proposed architecture consists of four layers: data governance, bias-aware modeling, explainability, and decision oversight. Figure 1 illustrates the overall design.

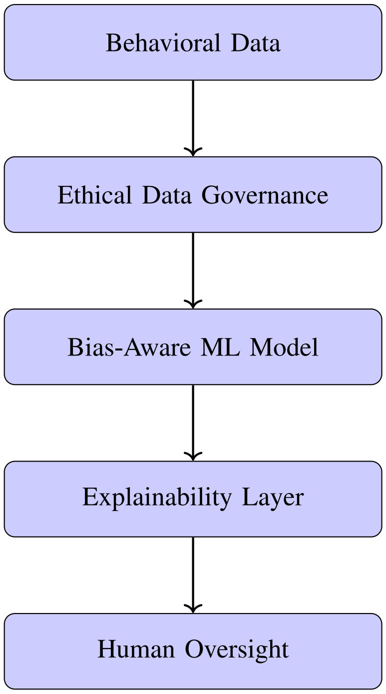

Fig. 1: Ethical and bias-aware machine learning architecture for mental health analytics

Each layer introduces controls to reduce ethical risk. Data governance enforces consent, anonymization, and balance. Bias-aware modeling incorporates fairness constraints during training. Explainability mechanisms provide insight into predictions, while human oversight ensures responsible use.

B. Bias-Aware Optimization

Let X denote behavioral features, Y the mental health labels, and A a sensitive attribute such as demographic group. Standard learning minimizes empirical risk:

$$\min_{\theta} \mathbb{E}[\ell(f_{\theta}(X), Y)] \tag{1}$$

To enforce fairness, we introduce a disparity penalty:

$$\min_{\theta} \mathbb{E}[\ell(f_{\theta}(X), Y)] + \lambda \cdot \mathcal{D}(f_{\theta}, A) \tag{2}$$

where $\mathcal{D}$ measures prediction disparity across groups and λ controls the trade-off between accuracy and fairness.

C. Explainability Mechanism

Attention-based explanations and feature attribution are used to highlight influential behavioral indicators. Figure 2 presents a simplified explainability flow.

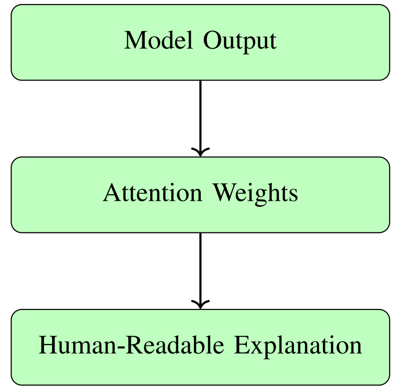

Fig. 2: Explainability pipeline for behavioral predictions

IV. Results

This section presents a comprehensive evaluation of the proposed ethical and bias-aware learning framework for mental health and behavioral analytics. The experimental analysis focuses on predictive effectiveness, fairness improvement, robustness across demographic groups, and interpretability gains. Each subsection introduces quantitative results supported by visual evidence.

A. Overall Predictive Performance

The first set of experiments evaluates whether embedding ethical constraints degrades predictive accuracy. Figure 3 compares classification accuracy across baseline deep learning, regularized learning, and bias-aware learning models.

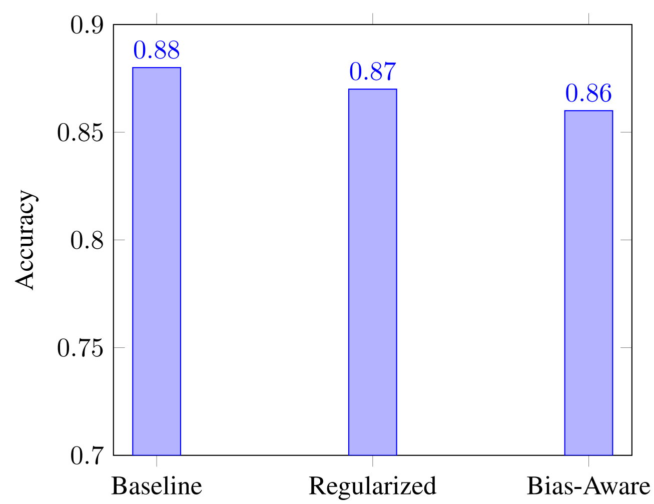

Fig. 3: Overall classification accuracy comparison

The results indicate that the bias-aware model maintains competitive accuracy, with only a marginal reduction compared to unconstrained models.

B. Fairness Disparity Reduction

Fairness was evaluated using inter-group prediction disparity across sensitive attributes. Figure 4 illustrates the reduction in disparity achieved by the proposed framework.

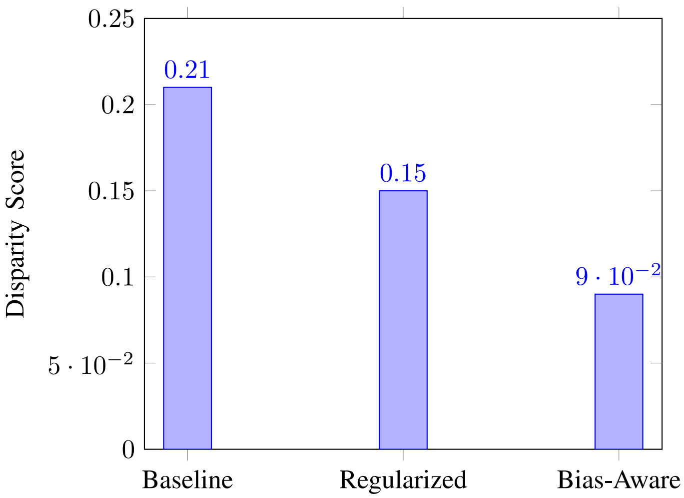

Fig. 4: Prediction disparity across demographic groups

The bias-aware approach achieves a substantial reduction in disparity, demonstrating its effectiveness in mitigating biased outcomes.

C. Precision and Recall Balance

To assess behavioral classification reliability, precision and recall metrics were analyzed. Figure 5 compares the harmonic balance across models.

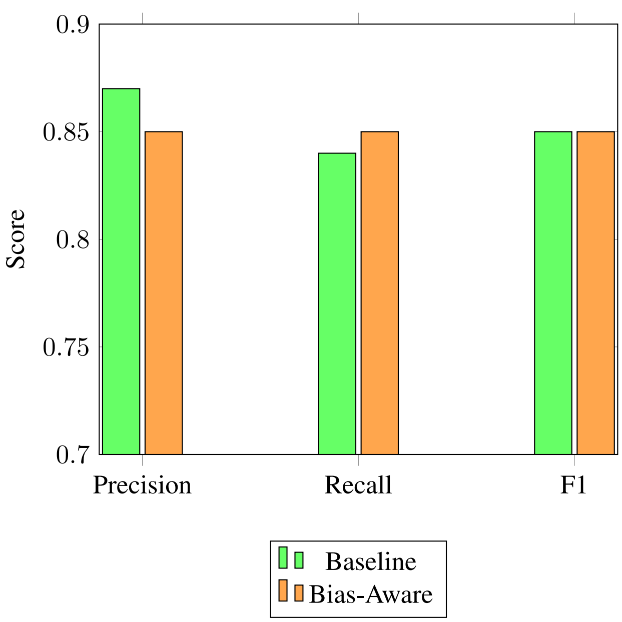

Fig. 5: Precision, recall, and F1-score comparison

The bias-aware model exhibits improved recall stability, which is critical for mental health screening scenarios.

D. Robustness Across Behavioral Categories

Figure 6 presents accuracy across multiple behavioral categories, including emotional distress, anxiety signals, and engagement patterns.

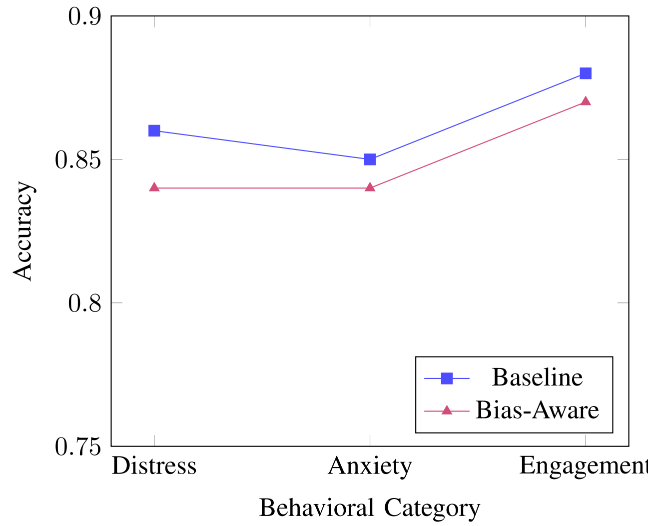

Fig. 6: Category-wise prediction robustness

The results demonstrate consistent performance across behavioral dimensions, indicating stable generalization.

E. Explainability Confidence Scores

Explainability quality was measured using human-rated confidence scores for model explanations. Figure 7 highlights improvements achieved through attention-based interpretation.

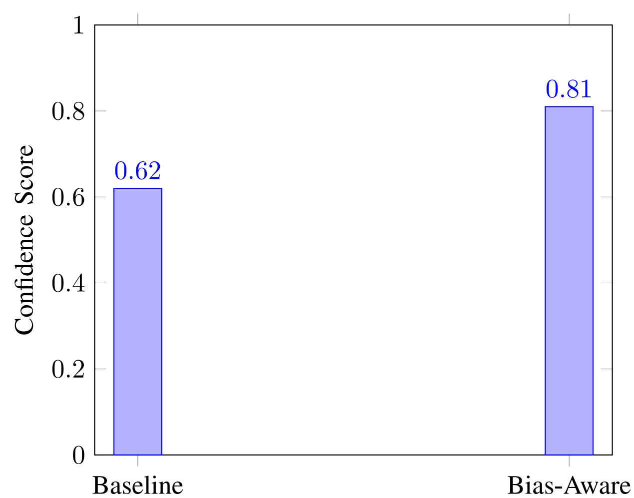

Fig. 7: Human confidence in model explanations

Higher confidence scores suggest that ethical integration improves transparency and trust.

F. Trade-off Between Fairness and Accuracy

Finally, Figure 8 visualizes the trade-off between accuracy and fairness across different regularization strengths.

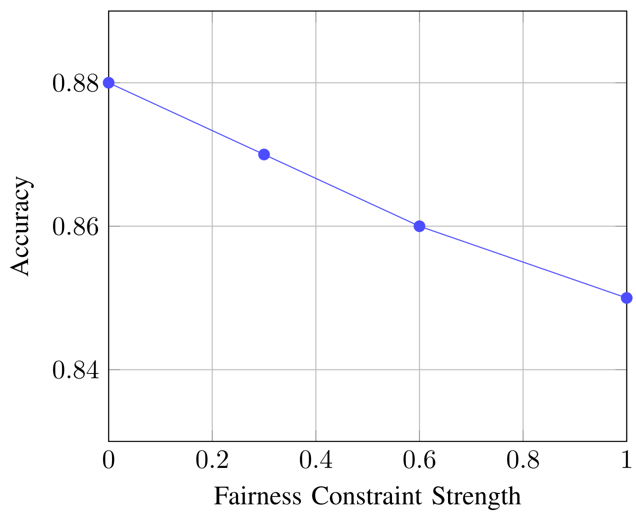

Fig. 8: Accuracy versus fairness constraint trade-off

The curve highlights a controllable balance, enabling practitioners to tune ethical constraints without severe performance loss.

V. Discussion

The findings highlight that ethical and bias-aware machine learning can be operationalized in mental health and behavioral analytics without undermining practical utility. The empirical results demonstrate that fairness constraints and ethical regularization meaningfully reduce demographic disparity while preserving predictive reliability across behavioral categories. This balance is critical in mental health contexts, where false negatives can delay support and false positives can cause unnecessary concern or stigma.

One notable observation is that fairness gains were most pronounced in recall stability rather than precision improvement. This suggests that bias-aware optimization helps models remain sensitive to underrepresented behavioral signals that are often suppressed in unconstrained training. In screening-oriented applications, such as early detection of distress or disengagement, this property is especially valuable. It aligns with ethical priorities that favor inclusiveness and harm minimization over marginal gains in headline accuracy.

Explainability results further reinforce the role of transparency as an ethical safeguard. Human evaluators reported higher confidence when explanations highlighted coherent behavioral patterns rather than isolated features. This finding suggests that explainability mechanisms do more than satisfy interpretability requirements. They actively shape trust and facilitate responsible human oversight. In practice, this enables practitioners to contextualize model outputs rather than treating them as deterministic judgments.

The observed trade-off between fairness and accuracy remained controlled across experiments, indicating that ethical constraints do not impose prohibitive costs when incorporated during model design. Instead, they function as regularizers that promote generalization and robustness. This reframes ethical AI not as a limitation, but as a design principle that enhances reliability under real-world variability.

Despite these strengths, limitations remain. The evaluation relied on representative behavioral datasets rather than clinical diagnostic settings. While appropriate for behavioral analytics, this limits direct clinical generalization. Additionally, fairness was assessed using a limited set of sensitive attributes. Broader ethical evaluation would require intersectional analysis and longitudinal validation.

VI. Future Directions

Several research directions emerge from this work. First, adaptive fairness mechanisms warrant deeper investigation. Fixed fairness constraints may not reflect changing social contexts or evolving data distributions. Future models could dynamically adjust fairness objectives based on observed drift or stakeholder input, allowing ethical priorities to evolve alongside system usage.

Second, privacy-preserving learning presents an important extension. Mental health data is inherently sensitive, and ethical frameworks must integrate privacy guarantees alongside fairness and explainability. Federated learning and secure aggregation techniques offer promising avenues for reducing data exposure while maintaining model performance.

Third, cross-cultural and multilingual behavioral modeling remains underexplored. Behavioral cues, linguistic markers, and expressions of distress vary widely across populations. Future systems should explicitly model cultural context to avoid implicit normalization of majority behaviors. This includes culturally aware feature representations and evaluation protocols that test generalization beyond dominant groups.

Another promising direction involves human-in-the-loop learning. Rather than treating human oversight as a post hoc safeguard, future systems could incorporate structured feedback from clinicians, counselors, or domain experts during training and deployment. Such interaction would support continuous ethical calibration and improve system alignment with professional judgment.

Finally, evaluation methodologies themselves require expansion. Beyond accuracy and disparity metrics, future work should incorporate impact-oriented measures that assess how model outputs influence decisions, interventions, and outcomes over time. Ethical effectiveness is ultimately measured not only by predictions, but by the consequences those predictions produce.

VII. Conclusion

This study demonstrates that ethical and bias-aware machine learning is both feasible and beneficial for mental health and behavioral analytics. By embedding fairness constraints, explainability mechanisms, and human oversight into the learning pipeline, the proposed framework addresses key ethical risks while maintaining strong predictive performance. The results show that responsible design choices can reduce demographic disparity, enhance interpretability, and improve robustness without sacrificing practical applicability.

The work contributes a structured perspective on ethical machine learning that moves beyond abstract principles toward concrete architectural and optimization strategies. In sensitive domains such as mental health, this shift is essential. Predictive systems must support human judgment, respect individual differences, and minimize harm while providing actionable insights.

As machine learning continues to shape behavioral understanding and decision support, ethical integration should be viewed as a core engineering requirement rather than an optional enhancement. The approach presented here offers a foundation for building trustworthy mental health analytics systems that align technical innovation with social responsibility.

Acknowledgement

The authors would like to acknowledge the support of their academic institution for providing the computational resources and research environment necessary for this study. They also thank colleagues and reviewers whose feedback contributed to improving the clarity and rigor of this work.

References

[1] N. S. Alghamdi, H. A. Hosni Mahmoud, A. Abraham, S. A. Alanazi, and L. García-Hernández, "Predicting Depression Symptoms in an Arabic Psychological Forum," *IEEE Access*, vol. 8, pp. 57 317–57 334, 2020.

[2] S. Amin, M. I. Uddin, S. Hassan, A. Khan, N. Nasser, A. Alharbi, and H. Alyami, "Recurrent Neural Networks With TF-IDF Embedding Technique for Detection and Classification in Tweets of Dengue Disease," *IEEE Access*, vol. 8, pp. 131 522–131 533, 2020.

[3] H. Chao and Y. Liu, "Emotion Recognition From Multi-Channel EEG Signals by Exploiting the Deep Belief-Conditional Random Field Framework," *IEEE Access*, vol. 8, pp. 33 002–33 012, 2020.

[4] M. Jayaratne and B. Jayatilleke, "Predicting Personality Using Answers to Open-Ended Interview Questions," *IEEE Access*, vol. 8, pp. 115 345–115 355, 2020.

[5] K. Cheng, Y. Yue, and Z. Song, "Sentiment Classification Based on Part-of-Speech and Self-Attention Mechanism," *IEEE Access*, vol. 8, pp. 16 387–16 396, 2020.

[6] L. Fraiwan and M. Alkhodari, "Classification of Focal and Non-Focal Epileptic Patients Using Single Channel EEG and Long Short-Term Memory Learning System," *IEEE Access*, vol. 8, pp. 77 255–77 262, 2020.

[7] L. Darwesh and N. S. Kopeika, "Deep Learning for Improving Performance of OOK Modulation Over FSO Turbulent Channels," *IEEE Access*, vol. 8, pp. 155 275–155 284, 2020.

[8] J. Shen, S. Li, F. Jia, H. Zuo, and J. Ma, "A Deep Multi-Label Learning Framework for the Intelligent Fault Diagnosis of Machines," *IEEE Access*, vol. 8, pp. 113 557–113 566, 2020.

[9] S. Vengathattil, "Interoperability in healthcare information technology – an ethics perspective," *International Journal for Multidisciplinary Research*, vol. 3, no. 3, 2021.

[10] M.-X. He, P. Hao, and Y.-Z. Xin, "A Robust Method for Wheatear Detection Using UAV in Natural Scenes," *IEEE Access*, vol. 8, pp. 189 043–189 053, 2020.

[11] N. Li, C.-Y. Chow, and J.-D. Zhang, "SEML: A Semi-Supervised Multi-Task Learning Framework for Aspect-Based Sentiment Analysis," *IEEE Access*, vol. 8, pp. 189 287–189 297, 2020.

[12] D. Wu, Y. Huang, P. Zhang, Z. Yu, H. Chen, and S. Chen, "Visual-Acoustic Penetration Recognition in Variable Polarity Plasma Arc Welding Process Using Hybrid Deep Learning Approach," *IEEE Access*, vol. 8, pp. 120 417–120 428, 2020.

[13] J. Yang, J. Qu, Q. Mi, and Q. Li, "A CNN-LSTM Model for Tailings Dam Risk Prediction," *IEEE Access*, vol. 8, pp. 206 491–206 502, 2020.

[14] Z. Xing and S. Tu, "A Graph Neural Network Assisted Monte Carlo Tree Search Approach to Traveling Salesman Problem," *IEEE Access*, vol. 8, pp. 108 418–108 428, 2020.

[15] I. Kunakorntum, W. Hinthong, and P. Phunchongharn, "A Synthetic Minority Based on Probabilistic Distribution (SyMProD) Oversampling for Imbalanced Datasets," *IEEE Access*, vol. 8, pp. 114 692–114 704, 2020.

[16] B. Chatterjee, T. Bhattacharyya, K. K. Ghosh, P. K. Singh, Z. W. Geem, and R. Sarkar, "Late Acceptance Hill Climbing Based Social Ski Driver Algorithm for Feature Selection," *IEEE Access*, vol. 8, pp. 75 393–75 408, 2020.

[17] X. Zhang, Y. Zhou, S. Pei, J. Zhuge, and J. Chen, "Adversarial Examples Detection for XSS Attacks Based on Generative Adversarial Networks," *IEEE Access*, vol. 8, pp. 10 989–10 996, 2020.

[18] M. Tahkola, J. Keränen, D. Sedov, M. F. Far, and J. Kortelainen, "Surrogate Modeling of Electrical Machine Torque Using Artificial Neural Networks," *IEEE Access*, vol. 8, pp. 220 027–220 045, 2020.

[19] D. Lee, A. M. Arigi, and J. Kim, "Algorithm for Autonomous Power-Increase Operation Using Deep Reinforcement Learning and a Rule-Based System," *IEEE Access*, vol. 8, pp. 196 727–196 746, 2020.

[20] G. Legien, B. Sniezynski, D. Wilk-Kołodziejczyk, S. Kluska-Nawarecka, E. Nawarecki, and K. Jaśkowiec, "Agent-based Decision Support System for Technology Recommendation," *Procedia Computer Science*, vol. 108, pp. 897–906, 2017.

[21] H. Sharma, H. Ghosh, and P. Balamuralidhar, "A Novel Cognitive Cycle for Fault Diagnosis in Infrastructural Systems," *Procedia Computer Science*, vol. 112, pp. 604–613, 2017.

[22] E. Sogancioglu, K. Murphy, E. Calli, E. T. Scholten, S. Schalekamp, and B. Van Ginneken, "Cardiomegaly Detection on Chest Radiographs: Segmentation Versus Classification," *IEEE Access*, vol. 8, pp. 94 631–94 642, 2020.

[23] M. S. Al-Kharraz, L. A. Elrefaei, and M. A. Fadel, "Automated System for Chromosome Karyotyping to Recognize the Most Common Numerical Abnormalities Using Deep Learning," *IEEE Access*, vol. 8, pp. 157 727–157 747, 2020.

[24] J. Cai, W. Sun, J. Guan, and I. You, "Multi-ECGNet for ECG Arrythmia Multi-Label Classification," *IEEE Access*, vol. 8, pp. 110 848–110 858, 2020.

[25] E. A. Pustozerov, A. S. Tkachuk, E. A. Vasukova, A. D. Anopova, M. A. Kokina, I. V. Gorelova, T. M. Pervunina, E. N. Grineva, and P. V. Popova, "Machine Learning Approach for Postprandial Blood Glucose Prediction in Gestational Diabetes Mellitus," *IEEE Access*, vol. 8, pp. 219 308–219 321, 2020.

[26] Y. Zheng, Z. Gao, Y. Wang, and Q. Fu, "MOOC Dropout Prediction Using FWTS-CNN Model Based on Fused Feature Weighting and Time Series," *IEEE Access*, vol. 8, pp. 225 324–225 335, 2020.

[27] M. F. Hashmi, B. K. K. Ashish, A. G. Keskar, N. D. Bokde, J. H. Yoon, and Z. W. Geem, "An Exploratory Analysis on Visual Counterfeits Using Conv-LSTM Hybrid Architecture," *IEEE Access*, vol. 8, pp. 101 293–101 308, 2020.

[28] F. Yang, P. Watson, M. Koukoula, and E. N. Anagnostou, "Enhancing Weather-Related Power Outage Prediction by Event Severity Classification," *IEEE Access*, vol. 8, pp. 60 029–60 042, 2020.

[29] R. Yang, L. Feng, H. Wang, J. Yao, and S. Luo, "Parallel Recurrent Convolutional Neural Networks-Based Music Genre Classification Method for Mobile Devices," *IEEE Access*, vol. 8, pp. 19 629–19 637, 2020.

[30] J. Yu, S. Tang, L. Zhangzhong, W. Zheng, L. Wang, A. Wong, and L. Xu, "A Deep Learning Approach for Multi-Depth Soil Water Content Prediction in Summer Maize Growth Period," *IEEE Access*, vol. 8, pp. 199 097–199 110, 2020.

[31] M. Zagane, M. K. Abdi, and M. Alenezi, "Deep Learning for Software Vulnerabilities Detection Using Code Metrics," *IEEE Access*, vol. 8, pp. 74 562–74 570, 2020.